

22883

1 **PATENT APPLICATION**

2

3 This application is submitted in the name of the following inventors:

4

5 *Inventor* *Citizenship* *Residence City and State*

6 Scott LOUGHMILLER United States Mountain View, California

7 Mike OLSON United States Sunnyvale, California

8 Jeff READY United States San Jose, California

9 Ehren MAEDGE United States Los Altos, California

10 Phil WHITE United States Santa Clara, California

11 Jason COLLIER United States San Jose, California

12

13 The assignee is *Corvigo*, a corporation having an office in Los Altos, Cali-
14 fornia.

15

16 **TITLE OF THE INVENTION**

17

18 Dynamic Message Filtering

19

1 **BACKGROUND OF THE INVENTION**

2

3 1. *Field of the Invention*

4

5 The invention relates to dynamic message filtering, such as for example
6 filtering incoming messages in response to their content; in one embodiment, messages
7 can be delivered, or other action taken, in response to a result of dynamic filtering.

8

9 2. *Related Art*

10

11 In computer communication networks, it is common to send and receive
12 messages among users, such as for the purpose of correspondence, distributing infor-
13 mation, and responding to requests. One method for doing so is electronic mail, also
14 known as email. One problem that has arisen in the art is that some messages are un-
15 wanted. Moreover, it has become common for advertisers and other message senders
16 to collect relatively large numbers of email addresses, and to send unsolicited advertis-
17 ing in bulk to recipients at those email addresses. When the number of such unsolicited
18 bulk email messages is relatively large, it can take substantial time and effort for recipi-
19 ents to delete them. There is also the possibility that the recipient will miss a relatively
20 important message due to the relatively large number of unimportant messages accu-
21 mulated in their email inbox. Such unsolicited bulk email messages are often known by

1 the colloquial term “spam,” and senders of such messages are often known as “spam-
2 mers.”

3

4 A first known method for detecting spam includes so-called “whitelists”
5 and “blacklists,” in which the sender of each message is identified by the filter as
6 known to be “good” (a sender who is not a spammer), or “bad” (a sender who is known
7 to be a spammer). While these methods generally achieve the goal of filtering mes-
8 sages, they are subject to the drawback that the user is involved in managing the
9 whitelist or blacklist, and the further drawback that spammers often choose new,
10 unique, sending addresses from which to send new spam.

11

12 A second known method for detecting spam includes attempting to
13 evaluate from the content of the message whether it is spam or not. Known evaluation
14 techniques include (a) searching the message for known keywords that are typically in-
15 dicative of spam, such as words identifying known products popularly promoted by
16 spammers, and (b) evaluating the message by comparing the number of such “bad”
17 keywords with probable “good” keywords, such as words relatively unlikely to be used
18 in a spam message. One example of the latter method is the Bayesian filter proposed by
19 Paul Graham, “A Plan for Spam,” and performed by some implementations of the
20 “Mozilla” email client. While these methods generally achieve the goal of filtering mes-
21 sages, they are subject to the drawback that the user must train the implementation to
22 recognize the “bad” keywords and “good” keywords particular to the type of message

1 that user typically receives, and the further drawback that spammers often choose, new,
2 unique, products to promote or words (often misspellings) with which to identify them.

3
4 Accordingly, it would be advantageous to provide an improved technique
5 for dynamic message filtering.
6

7 SUMMARY OF THE INVENTION 8

9 The invention provides a method and system capable of dynamically fil-
10 tering incoming messages, with the effect of classifying those messages into one of at
11 least three categories: good messages, bulk periodicals, and spam. The intermediate
12 category of "bulk periodicals" is reserved for messages that are clearly not directed to
13 the individual recipient, but which the recipient might wish to review anyway, such as
14 for example information relating to updates of products the user is already using, or in-
15 formation relating to products or services the user is explicitly interested in.
16

17 In a first aspect, the system includes an identification engine that classifies
18 messages based on a measured intent of each message. In one embodiment, the engine
19 includes a regular expression recognizer and a set of artificial neural networks pre-
20 trained to classify messages. The regular expression recognizer is suitable for detecting
21 misspelled words, likely spam phrases composed of otherwise innocent words (such as
22 for example "MAKE MONEY FAST"), and other common attempts by spammers to

1 evade detection by known keywords that are typically indicative of spam. The artificial
2 neural networks divide messages into “likely good” and “likely spam,” and with that
3 information, operate at a more detailed and discriminating level to distinguish among
4 good messages, bulk periodicals, and spam. Messages initially considered “likely
5 good” might be ultimately identified as good messages or as bulk periodicals. Simi-
6 larly, messages initially considered “likely spam” might be ultimately identified as bulk
7 periodicals or as spam. This aspect accounts for the fuzziness in determination, and re-
8 duces the number of messages erroneously identified as spam by identifying a signifi-
9 cant number of them as bulk periodicals, which are considered relatively less pernicious
10 by the user.

11

12 In a second aspect, the system includes a dynamic whitelist and blacklist,
13 into which sending addresses are collected when the number of messages from those
14 sending addresses indicates that the sender is likely good or likely a spammer. In one
15 embodiment, any sender for whom at least a threshold number (preferably four) of
16 messages pass as good messages is automatically added to the whitelist of known good
17 senders, so that messages from those senders need not be checked as thoroughly as
18 from other senders.

19

20 In a third aspect, the system includes a set of regular expressions whose
21 detection is input to the artificial neural networks, in one embodiment selected before
22 installation, with the effects that the artificial neural networks can be trained more rap-

1 idly, and respond more rapidly and accurately to changes in the typical email received
2 by the user. In one embodiment, a subset of the 2,000 most useful regular expressions
3 (identifying words or phrases) is selected using a genetic algorithm, out of the possibly
4 70,000 most common English words and phrases that might be used. This also has the
5 effect that the artificial neural networks can be made smaller (that is, with fewer input
6 nodes and fewer hidden nodes), and are capable of being executed directly in relatively
7 less main memory, with the effect that such execution is relatively faster.

8

9 The invention is not restricted to email messages, but is also applicable to
10 other types of messages or data, such as for example web pages or web page caching,
11 "pop-up" advertising, and web page JavaScript, "instant messages," message protocols
12 using HTTP tunneling, as well as to other types of filtering, such as for example auto-
13 matic routing of email to appropriate recipients, automatic prioritization for review (or
14 for forwarding to a pager or wireless email inbox), automatic collection of groups of
15 messages into message digests, automatic flagging of messages for secondary review or
16 for legal liability, and automatic detecting of outgoing messages for virus content.

17

18 **BRIEF DESCRIPTION OF THE DRAWINGS**

19

20 Figure 1 shows a block diagram of a generalized system for dynamic mes-
21 sage filtering.

22

1 Figure 2 shows a block diagram of a system for dynamic message filtering,
2 in an embodiment disposed behind a firewall.

3
4
5
6

7 Figure 4 shows a block diagram of one embodiment of an identification
8 engine.

9

10 Figure 5 shows a flow diagram of a method for dynamic message filtering.

11

12 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT**

13

14 In the description herein, a preferred embodiment of the invention is de-
15 scribed, including preferred process steps and data structures. Those skilled in the art
16 would realize, after perusal of this application, that embodiments of the invention
17 might be implemented using a variety of other techniques not specifically described,
18 without undue experimentation or further invention, and that such other techniques
19 would be within the scope and spirit of the invention.

20

21 *Lexicography*

22

1 The general meaning of each of these following terms is intended to be il-
2 lustrative and in no way limiting.

3

4 • The terms “email,” “message” and the phrases “electronic mail” and “electronic
5 message” generally describe any technique by which information is carried (or
6 attempted to be carried) from a sender to a recipient, whether that technique is
7 asynchronous or synchronous, circuit switched or packet switched, human
8 readable or not, intended for individual recipients or not, or broadcast or
9 multicast or unicast or not. As used herein, the term “message” is intentionally
10 broad. As described herein, there is no particular requirement that messages
11 must be pure or primarily text.

12

13 • The phrases “unwanted email,” “unwanted messages” and the term “spam”
14 generally describe any message not specifically desired by, or offensive to, or
15 wasting the time of, its actual or potential recipient. As used herein, the term
16 “spam” is intentionally broad, even though it might be typically applied to
17 messages that are unsolicited, sent in bulk, and often involve advertising.

18

19 • The phrase “bulk periodicals” generally describes, when applied to messages,
20 messages that are clearly not directed to the individual recipient, but which the
21 recipient might wish to review anyway, such as for example information relating

1 to updates of products the user is already using, or information relating to prod-
2 ucts or services the user is explicitly interested in.

3

- 4 • The phrase “relatively important message” generally describes any message not
5 considered “spam,” and not considered a desirable bulk message, such as for
6 example a periodical or an advertisement of interest to the recipient. This has the
7 effect that the invention provides a method and system capable of dynamically
8 filtering incoming messages, with the effect of classifying those messages into
9 one of at least three categories: good messages, bulk periodicals, and spam. The
10 former are relatively important messages, the latter are relatively unwanted
11 messages, and the middle (bulk periodicals) are messages that might or might
12 not be desirable to the user.

13

- 14 • The terms “blocking” and “filtering” when applied to messages, generally
15 describe any technique by which those messages are marked for action on the
16 part of a user, such as for example to tag them with an evaluation of whether the
17 message is spam, to take action in response to that evaluation, or to order or
18 otherwise prioritize those messages in response to that evaluation. For example,
19 not intended to be limiting in any way, “taking action” in response to that
20 evaluation might include forwarding part or all of a message to a wireless
21 recipient, copying the message to a more permanent log, redistributing the
22 message to another user, or reporting the sender of the message to an anti-spam

1 enforcer (such as for example the spammer's ISP, a common blacklist of
2 spammer sending addresses, or a government agency).

3

- 4 • The phrase "false positives", when applied to messages, generally describe any
5 messages misidentified as not being relatively important messages, when in fact
6 they are.

7

8 Other and further applications of the invention, including extensions of
9 these terms and concepts, would be clear to those of ordinary skill in the art after pur-
10 chasing this application. These other and further applications are part of the scope and
11 spirit of the invention, and would be clear to those of ordinary skill in the art without
12 further invention or undue experimentation.

13

14 The scope and spirit of the invention is not limited to any of these defini-
15 tions, or to specific examples mentioned therein, but is intended to include the most
16 general concepts embodied by these and other terms.

17

18 *Generalized System Elements*

19

20 Figure 1 shows a block diagram of a generalized system for dynamic mes-
21 sage filtering.

1 In a generalized conception of an embodiment of the invention, a system
2 100 is coupled to an outside network 110, and includes an appliance level 120, a server
3 level 130, and a client level 140.

4

5 The outside network 110 might include any technique for sending or re-
6 ceiving messages, including incoming mail traffic 111 such as email and other messages.
7 In one embodiment, the outside network 110 includes an Internet, such as for example
8 an ISP coupled to an Internet backbone network. However, in the context of the inven-
9 tion, there is no particular requirement that the outside network 110 involves any par-
10 ticular type of communication system. In alternative embodiments, the outside net-
11 work 110 may include an intranet, extranet, VPN, an ATM network, a private or public
12 switched network such as for example a PSTN, or some combination or conjunction
13 thereof.

14

15 In one embodiment, the appliance level 120 includes an entry point 125 to
16 an enterprise network, possibly involving a firewall, a router or gateway router, or a
17 software construct such as a VPN (virtual private network) disposed within a more in-
18 clusive communication network. The appliance level 120 includes a spam filter 121
19 coupled to the entry point to the enterprise network, and also coupled to the rest of the
20 enterprise network. For one example, not intended to be limiting in any way, the spam
21 filter 121 might be coupled to the firewall at a first port 123, and coupled to the enter-
22 prise network (such as a network including a LAN, WAN, VLAN, VPN, or the like) at a

1 second port 124. In one embodiment, the spam filter 121 includes at least some storage
2 122, such as for example a database or other storage, at which the spam filter 121 might
3 maintain any junk mail (spam) blocked, retained, or withheld by the spam filter 121.

4

5 In one embodiment, the server level 130 includes a set of server elements.
6 For example, not intended to be limiting in any way, the server level 130 might include
7 at least one of: a mail server, a web server, a database server, a server for network-

8 attached storage, or a print server. In one embodiment, the server level 130 includes at
9 least one mail server 131, which is coupled to the spam filter 121 at a first port 132, and
10 coupled to the client level 140 at a second port 133. For example, not intended to be
11 limiting in any way, the mail server 131 might be coupled to a set of workstations at
12 which users receive and manipulate email messages.

13

14 In one embodiment, the client level 140 includes a set of client devices.
15 For example, not intended to be limiting in any way, the client level 140 might include a

16 set of workstations, printers, wireless communication devices, or handheld devices such
17 as for example "Blackberry" or "Palm Pilot" devices or PDA's (personal digital assis-
18 tants) or personal organizers. In one embodiment, the client level 140 includes at least
19 one recipient mailbox 141. The recipient mailbox 141 includes at least two regions, a
20 "good messages" mailbox section 142 and a "bulk periodicals" mailbox section 143.

21

1 As described below, the spam filter 121 receives at least some of the in-
2 coming mail traffic 111 from the outside network 110, and classifies messages from that
3 incoming mail traffic 111 into a set of classes. In one embodiment, this set of classes in-
4 cludes "good messages," "bulk periodicals," and "junk email." The spam filter 121 re-
5 tains those messages classified as junk email at the appliance level 120 in storage 122,
6 while allowing those messages classified as good messages or bulk periodicals, suitably
7 marked, to be forwarded to the server level 130. At the server level 130, the mail server
8 131 receives the good messages or bulk periodicals, suitably marked, delivers the good
9 messages to the "good messages" mailbox section 142, and delivers the bulk periodicals
10 to the "bulk periodicals" mailbox section 143.

11

12 *System Elements (Behind a Firewall)*

13

14 Figure 2 shows a block diagram of a system for dynamic message filtering,
15 in an embodiment disposed behind a firewall.

16

17 An embodiment of a system 200 includes an outside network 210 (similar
18 to the outside network 110 of figure 1), a firewall 220 (similar to the firewall, router or
19 gateway router described with regard to the appliance level 120 of figure 1), a mail
20 server 230 (similar to the mail server 131 of figure 1), an administrator web client 241, an
21 end-user web client 242, and a spam filter 250 (similar to the spam filter 121 of figure 1).

22

1 Similar to the outside network 110 of figure 1, the outside network 210
2 might include any technique for sending or receiving messages. In one embodiment,
3 the outside network 210 includes an Internet, such as for example an ISP coupled to an
4 Internet backbone network. However, in the context of the invention, there is no par-
5 ticular requirement that the outside network 210 involves any particular type of com-
6 munication system. In alternative embodiments, the outside network 210 may include
7 an intranet, extranet, VPN, an ATM network, a private or public switched network such
8 as for example a PSTN, or some combination or conjunction thereof.

9

10 Similar to the firewall described with regard to the appliance level 120 of
11 figure 1, the firewall 220 is capable of receiving mail traffic such as email and other mes-
12 sages 221 from the outside network 210, examining those messages 221 to determine if
13 they should be blocked or not (in response to a set of firewall rules maintained by the
14 firewall 220), and sending those messages 221 (if not blocked) to the spam filter 250.

15

16 Similar to the mail server 131 of figure 1, the mail server 230 is capable of
17 receiving messages 221 from the spam filter 250 and forwarding those messages to end-
18 user recipients in response to their contents. For some examples, not intended to be
19 limiting in any way, the mail server 230 might include a known mail server using a
20 known mail transfer protocol, such as for example IMAP, POP, or SMTP.

21

1 The administrator web client 241 includes a processor, program and data
2 memory, and input and output devices, such as for example configured as a desktop
3 workstation, a notebook computer, a "Blackberry" or "Palm Pilot" or other handheld
4 computing device, or other device. The administrator web client 241 is capable of
5 communicating with the spam filter 250, with the effect than an administrator is capable
6 of reviewing, editing, or deleting configuration information maintained by the spam
7 filter 250 for general use.

8

9 The end-user web client 242 includes a processor, program and data
10 memory, and input and output devices, such as for example configured as a desktop
11 workstation, a notebook computer, a "Blackberry" or "Palm Pilot" or other handheld
12 computing device, or other device. The end-user web client 241 is capable of communi-
13 cating with the spam filter 250, with the effect than an end-user is capable of reviewing,
14 editing, or deleting configuration information maintained by the spam filter 250 for use
15 involving that end-user.

16

17 *Spam Filter*

18

19 Similar to the spam filter 121 of figure 1, the spam filter 250 is capable of
20 receiving the messages 221 allowed through by the firewall 220, examining those mes-
21 sages 221 to determine if they should be treated as good messages, bulk advertising, or
22 spam, and taking one or more actions with regard to those messages 221 in response to

1 a result of that determination. Those one or more actions might include (a) tagging the
2 message 221 appropriately before forwarding it to the mail server 230 for delivery, (b)
3 delaying, deleting, quarantining, or otherwise treating the message 221, (c) forwarding
4 the message 221 to users or other entities besides the end-user to whom the message 221
5 was originally addressed, and (d) updating its own state with the effect that the spam
6 filter 250 is better able to discern which messages 221 are good messages, bulk adver-
7 tising, or spam.

8

9 In one embodiment, the spam filter 250 includes a mail transfer agent 251,
10 a database 252, an identification engine 253, an administration interface element 254, an
11 end-user interface element 255, a web server 256, a web CGI layer 257, an operating
12 system layer 258, and a hardware platform 259.

13

14 In one embodiment, the mail transfer agent 251 includes a PostFix Mail
15 Transfer Agent, such as for example a version 1.1.11 (available from IBM), modified to
16 include capabilities and functions as described herein. The mail transfer agent 251
17 could also use or instead include Sendmail.

18

19 The mail transfer agent 251 is capable of transferring messages 221 among
20 or between devices, using the general schema that user senders (using user agents) send
21 messages 221 to the mail transfer agent 251, which sends the message 221 to one or
22 more other mail transfer agents, each of which delivers the message 221 to one or more

1 end-user recipients. In one embodiment, the mail transfer agent 251 is modified with
2 the effect that it communicates with the database 252 and the identification engine 253
3 to examine and classify messages 221.

4

5 In one embodiment, the database is used to store user and administrative
6 settings, as well as statistics and email logging and reporting. Messages that are identi-
7 fied as spam can also be stored in a file system for later retrieval if a user determines
8 that the messages are not actually spam (that is, they were misidentified “false posi-
9 tives”). In alternative embodiments, messages that are identified as spam can also be
10 stored in the database 252 for later retrieval under similar conditions.

11

12 In one embodiment, the identification engine 253 includes a Corvigo (pro-
13 prietary) filtering engine, such as for example version 2.4 thereof. In this embodiment,
14 the filtering engine uses a combination of artificial intelligence techniques, for example
15 including natural language processing, to determine the intent of messages. Filtering
16 can then be performed based on determined intent.

17

18 In one embodiment, the administration interface element 254 includes an
19 interface. The administrator interface element allows an administrator to configure, run
20 and maintain the spam filter 250. The administration interface element 254 might be
21 built using known tools, such as for example HTML (hypertext markup language) and
22 PHP.

1
2 In one embodiment, the end-user interface element 255 includes a user
3 interface. The user interface element allows users to perform one or more of the fol-
4 lowing operations: modifying their spam filtering preferences, viewing a report of mail
5 messages and how that mail was categorized and handled, and allowing the users to
6 find and retrieve "false positives" (that is, good messages mistakenly identified as bulk
7 periodicals or spam). The end-user interface element 255 might be built using known
8 tools, such as for example HTML and PHP.

9
10 In one embodiment, the web server 256 includes an open-source server
11 such as the Apache HTTP Server license 1.0. In this embodiment, the web server pro-
12 vides functions and interfaces used to generate a web CGI layer such as web CGI layer
13 257.

14
15 In one embodiment, the web CGI layer 257 includes a PHP (Hypertext
16 Preprocessor) based interface that allows users and administrators to interact with spam
17 filter 250 over an HTML-enabled network such as the World Wide Web.

18
19 In a first preferred embodiment, the hardware platform 259 and the oper-
20 ating system layer 258 include an Intel-architecture processor (or a functional equiva-
21 lent thereof) operating under control of a version of the Linux operating system (or a
22 similar operating system, such as a version of Unix or an operating system including

1 the Mach microkernel). In a second preferred embodiment, the hardware platform 259
2 and the operating system layer 258 include a Sun SPARC station processor (or a func-
3 tional equivalent thereof) operating under control of a version of the Solaris operating
4 system (or a similar operating system, such as a version of Unix).

5

6 As described below, in one embodiment the mail transfer agent 251 at the
7 spam filter 250 receives at least some of the messages 221 from the outside network 210,
8 such as for example possibly using the firewall 220. The messages 221 are routed to the
9 identification engine 253 for classification based on an intent of each message, as deter-
10 mined by that identification engine 253.

11

12 In one embodiment, the engine includes a regular expression recognizer
13 and a set of artificial neural networks pre-trained to classify messages. The regular ex-
14 pression recognizer is suitable for detecting misspelled words, likely spam phrases
15 composed of otherwise innocent words (such as for example "MAKE MONEY FAST"),
16 and other common attempts by spammers to evade detection by known keywords that
17 are typically indicative of spam. The artificial neural networks divide messages into
18 "likely good" and "likely spam," and with that information, operate at a more detailed
19 and discriminating level to distinguish among good messages, bulk periodicals, and
20 spam.

21

1 Messages initially considered “likely good” might be ultimately identified
2 as good messages or as bulk periodicals. Similarly, messages initially considered
3 “likely spam” might be ultimately identified as bulk periodicals or as spam.

4

5 In one embodiment, messages identified as spam are sent to a file system
6 for storage, in case they were misidentified (that is, they are “false positives”). In alter-
7 native embodiments, messages identified as spam may optionally be sent to the data-
8 base 252 for later identification or other processing. Good messages and bulk periodi-
9 cals, suitably marked, are sent from the mail transfer agent 251 to the mail server 230 for
10 delivery to mailbox sections for end-user recipients.

11

12 Different ways of handling the various types of messages also are possi-
13 ble. For one example, not intended to be limiting in any way, the spam filter 250 could
14 be configured to handle bulk messages as if they were spam.

15

16 *System Elements (Configured as a Server)*

17

18 After reading this application, those skilled in the art would recognize that
19 the system for dynamic message filtering may alternatively be configured for use in an
20 embodiment disposed as a server. In such embodiments, the server would be capable
21 of generally similar to a combination of the mail server 131 and the spam filter 121 of
22 Figure 1. This would have the effect that the server would be capable of receiving mes-

1 sages, filtering out spam and possibly bulk periodical messages, and forwarding good
2 messages to end-user recipients in response to their contents. In some embodiments,
3 not intended to be limiting in any way, the server might function using a known mail
4 transfer protocol, such as for example IMAP, POP, or SMTP.

5

6 *Identification Engine*

7

8 Figure 4 shows a block diagram of one embodiment of an identification
9 engine according to an aspect of the invention.

10

11 An embodiment of an identification engine 400, suitable for use as identi-
12 fication engines 250 or 350, includes a lexical analyzer 410, an input vector generator
13 420, rules 430, and a neural network hierarchy 440.

14

15 In one embodiment, the lexical analyzer 410 decodes and parses messages
16 according to internet standards. The message is broken down into a header section and
17 a body section. The header section is further decomposed into individual headers. The
18 body section is decoded if necessary and stripped of extraneous markup. In this em-
19 bodiment, multiple MIME parts and their subcomponents are handled.

20

21 In one embodiment, the input vector generator 420 includes a regular ex-
22 pression recognizer that uses the subject header and the text of the body to generate an

1 input vector suitable for use by the hierarchy of neural networks 440. This is accom-
2 plished by using a many-to-one map of words and phrases to input vector positions. In
3 one embodiment, the map includes a subset of the 2,000 most useful regular expressions
4 (identifying words or phrases) pre-selected using a genetic algorithm out of the possibly
5 70,000 most common English words that might be used.

6

7 In one embodiment, the input vector generator 420 scans the textual char-
8 acters of each message, and each word or phrase (up to four words long) that appears in
9 the message is checked for a value in the map. If an entry corresponding to the word or
10 phrase appears in the map, the appropriate value of the input vector is increased. Care
11 is taken so that words are recognized in the message even if the message is composed in
12 such a way as one might devise to circumvent the system.

13

14 In one embodiment, rules 430 apply basic rules to messages to possibly
15 determine its classification in an efficient manner. One possible set of such rules are
16 enumerated below. In this embodiment, once a classification has been positively de-
17 termined by any rule, no further processing occurs.

18

19 (1) A message is classified if the sender of the message exists in a list of
20 senders known to the recipient.

21

1 (2) A message is classified if the IP address of the client that sent the mes-
2 sage exists in a list of IP addresses known by the administrator. The list contains the
3 client IP address and the desired classification.

4

5 (3) A message is classified if the sender of the message exists in a list of
6 senders known to the administrator. The list contains the sender's address and the de-
7 sired classification.

8

9 (4) A message is classified if the domain part of the sender's address exists
10 in a list of domains known to the administrator. The list contains the sender's address'
11 domain and the desired classification.

12

13 (5) A message is classified as junk if the subject contains the string "adv"
14 (or another selected string) followed by a delimiter. Such strings are sometimes used by
15 advertisers to alert end users that the message is a commercial advertisement or other
16 bulk email.

17

18 (6) A message may be classified if it uses a character set that is not com-
19 monly used for U.S. English messages (or another selected character set or language).
20 The classification may occur in response to system configuration.

21

1 (7) A message is classified if its subject matches any search strings in a ta-

2 ble containing regular expressions, search strings, and associated classifications in re-

3 sponse thereto.

4

5 (8) A message is classified if its body matches any search strings in a table

6 containing search strings and classifications.

7

8 (9) A message is classified if any of its headers match any search strings in

9 a table containing search strings and classifications.

10

11 (10) A message is classified as junk if it contains code that would cause a

12 mail reader to automatically create a new web browser window.

13

14 (11) A message is classified if the recipient's address does not appear in

15 any of the standard message headers that contain lists of recipient addresses (such as

16 for example "To" and "CC" headers). The classification may occur based on system con-

17 figuration.

18

19 (12) A message may be classified as junk if the list of recipients as declared

20 by the message's headers contain mostly addresses beginning with the same letter.

21

1 (13) If insufficient input has been generated by the input vector generator
2 420 for the message, it will be classified as a "good message."

3
4 In other embodiments, different rules 430 may be used, including some,
5 all, or none of the foregoing examples.

6
7 With respect to the various lists used by the rules, in one embodiment
8 these lists are divided into "whitelists" that include identifiers for good messages and
9 "blacklists" that include identifiers for spam messages. Examples of identifiers include
10 but are not limited to a sender's name, address, domain name, or IP address.

11
12 In one embodiment, the whitelists and blacklists can be dynamically
13 maintained based on the classification of messages associated with those identifiers.
14 For example, any sender for whom at least a threshold number (preferably four) of
15 messages pass as good messages can be automatically added to the whitelist of known
16 good senders, so that messages from those senders need not be checked as thoroughly
17 as from other senders. Likewise, any sender for whom at least a threshold number
18 (preferably four) of messages are rejected as spam can be automatically added to the
19 blacklist of known spammers, so that messages from those senders need not be checked
20 as thoroughly as from other senders. These classifications can come from the overall
21 operation of the spam filter or spam filtering server, from user review of messages, or

1 from some other source. Different thresholds and techniques for dynamically updating
2 the lists also can be used.

3

4 If none of the rules 430 positively classify the message, in one embodiment,
5 the vector created by the input vector generator is processed by the hierarchy of
6 neural networks 440, further described with regard to figure 3. The neural networks
7 embody an artificial intelligence engine that filters the messages by looking at the intent
8 of the messages as indicated by the generated input vector. One embodiment of the
9 neural network hierarchy analyzes what words are used in a message, analyzes how
10 those words are used both independently and in relationship with each other, and (c)
11 considers a classification for the message based on this analysis and on knowledge of
12 other messages.

13

14 One embodiment of hierarchy 440 includes at least two neural networks.
15 The first neural network determines if a message is more likely legitimate or junk, di-
16 viding messages into "likely good" and "likely spam."

17

18 Based on the initial classification, a second neural network processes the
19 input vector to determine if the message is bulk mail. In this embodiment, there are
20 separate neural networks to classify bulk message from junk messages, and bulk mes-
21 sages from legitimate messages. Messages initially considered "likely good" might be
22 ultimately identified as good messages or as bulk periodicals. Similarly, messages ini-

1 tially considered "likely spam" might be ultimately identified as bulk periodicals or as
2 spam.

3

4 In alternative embodiments, the neural networks learn from messages that
5 are processed in order to adapt to evolving anti-filtering strategies employed by senders
6 of spam.

7

8 Different hierarchies with different numbers and purposes of neural net-
9 works can be used in other embodiments of the invention.

10

11 *Neural Networks*

12

13 Figure 3 shows a block diagram of one embodiment of a set of neural net-
14 works according to an aspect of the invention.

15

16 A system 300 of neural networks 310A, 310B, and 310C includes at least a
17 first neural network 310A, having a set of input nodes 311, a neural network body 312,
18 and an output node 313. In one embodiment, each one of the input nodes 311 is cou-
19 pled to a corresponding one regular expression recognizer 314.

20

21 A set of input words 315 from a message are coupled to the set of regular
22 expression recognizers 314. Each one of the regular expression recognizers 314 gener-

1 ates, in response to the set of input words 315, a value (0 or 1) representing the absence
2 or presence of an associated pattern, as represented by a regular expression. In one em-
3 bodiment, the regular expression recognizers 314 are pre-selected. However, in alter-
4 native embodiments, the regular expression recognizers 314 may be altered in response
5 to user feedback regarding whether a particular one or more messages are properly
6 identified.

7

8 After reading this application, those skilled in the art would recognize that
9 the regular expression recognizers 314 are not required to use regular expressions, or to
10 provide discrete values of 0 or 1 in response to the set of input words 315. For a first
11 example, not intended to be limiting in any way, the regular expression recognizers 314
12 might be replaced or assisted by other types of pattern matchers or machine learning
13 techniques. For a second example, not intended to be limiting in any way, the regular
14 expression recognizers 314 might use fuzzy logic or otherwise provide substantially
15 continuous values (or one of a set of discrete values) between 0 and 1.

16

17 The set of outputs from the regular expression recognizers 314 is coupled
18 to corresponding ones of the input nodes 311. Each one of the input nodes 311 is as-
19 signed a weighted value in response to a count of the number of regular expressions
20 identified by the regular expression recognizers 314 and associated with that input node
21 311. This has the effect that, if a particular regular expression is identified twice, the in-

1 put to that input node 311 will be in response to the value 2, rather than just 1 for identi-
2 fying that particular regular expression once.

3

4 In one embodiment, each individual number of identified regular expres-
5 sions is divided by the total number of identified regular expressions, with the effect
6 that the values coupled to the input nodes 311 are substantially normalized to a total of
7 1. For one example, not intended to be limiting in any way, if there are 4 input nodes
8 311A, 311B, 311C, and 311D, and the number of identified regular expressions for each
9 is 3, 3, 6, and 3 respectively (thus totaling 15), the normalized values will be 3/15, 3/15,
10 6/15, and 3/15 respectively (thus totaling 1.0).

11

12 In one embodiment, each substantially normalized value is adjusted to a
13 minimal non-zero value, with the effect that the values coupled to the input nodes 311
14 are only zero if the pre-normalized number of identified regular expressions was ex-
15 actly zero. In other cases, where the pre-normalized number of identified regular ex-
16 pressions was more than zero, but the normalized value was quite small (for example,
17 0.02), that normalized value is rounded up to a minimum quantum, preferably 0.1. In
18 alternative embodiments, the same process might be conducted for maximum values as
19 well.

20

21 The set of outputs from the input nodes 311 are coupled to a fully-
22 connected neural network 312, with, in one embodiment, thresholds and weights pre-

1 selected. However, in alternative embodiments, the thresholds or weights, or both, for
2 the neural network 312 may be adjusted in response to user feedback.

3

4 The outputs from the fully-connected neural network 312 are coupled to
5 an output node 313, with the effect that the output node 313 presents a value of between
6 0 and 1. A result of one of the neural networks 310A, 310B, or 310C, is responsive to a
7 threshold value associated with that neural network 310A, 310B, or 310C, such as for
8 example a preferred threshold value of 0.9. This has the effect that for the first neural
9 network 310A, if the threshold value is exceeded, the message is re-evaluated by the
10 second neural network 310B or if the threshold value is not exceeded, the message is re-
11 evaluated by the third neural network 310C. This also has the effect that for the second
12 neural network 310B, if the threshold value is exceeded, the message is determined to
13 be spam, or if the threshold value is not exceeded, the message is determined to be bulk
14 email of possible interest. Similarly, this also has the effect that for the third neural
15 network 310C, if the threshold value is exceeded, the message is determined to be bulk
16 email of possible interest, or if the threshold value is not exceeded, the message is de-
17 termined to be a "good message."

18

19 *Method of Operation*

20

21 Figure 5 shows a flow diagram of a method for dynamic message filtering.

22

1 Although described serially, the flow points and method steps of the
2 method 500 can be performed by separate elements in conjunction or in parallel,
3 whether asynchronously or synchronously, in a pipelined manner, or otherwise. In the
4 context of the invention, there is no particular requirement that the method must be
5 performed in the same order in which this description lists flow points or method steps,
6 except where explicitly so stated.

7

8 In a step 501, a message 321 is received.

9

10 In a step 502, the message 321 is sent to an identification engine, such as
11 for example the identification engine 400 in Figure 4.

12

13 In a step 503, a lexical analyzer such as for example the lexical analyzer
14 410 in Figure 4 decodes and parses the message according to known standards, such as
15 for example known Internet standards. The message is broken down into a header sec-
16 tion and a body section. The header section is further decomposed into individual
17 headers. The body section is decoded if necessary and stripped of extraneous markup.
18 In one embodiment, multiple MIME parts and their subcomponents are handled.

19

20 In a step 504, an input vector generator, such as the input vector generator
21 420 in Figure 4, recognizes words and expressions in the text of the header and body
22 sections. The input vector generator uses the recognized words and expressions to gen-

1 erate an input vector. The vector is generated using a many-to-one map of words and
2 phrases to input vector positions. In one embodiment, the map includes a subset of the
3 2,000 most useful regular expressions (identifying words or phrases) selected using a
4 genetic algorithm out of the possibly 70,000 most common English words that might be
5 used.

6

7 In one embodiment, the input vector generator scans the text of each mes-
8 sage, and each word or phrase (up to four words long) that appears in the message is
9 checked for a value in the map. If an entry corresponding to the word or phrase ap-
10 pears in the map, the appropriate value of the input vector is increased. Care is taken
11 so that words are recognized in the message even if the message is composed in such a
12 way as one might devise to circumvent the system.

13

14 In a step 505, rules such as for example rules 430 in Figure 4 are applied to
15 the message. In one embodiment, basic rules are applied to the message to possibly
16 determine its classification in an efficient manner.

17

18 If the rules successfully classify the message, flow proceeds from step 506
19 to step 507. In step 507, the spam filter or spam filtering server acts upon the message
20 based on the classification. One embodiment of possible classifications and actions is
21 explained in more detail below with respect to steps 511 to 513.

22

1 If the rules do not successfully classify the message, flow proceeds from
2 step 506 to step 508 to 510. In those steps, the message is analyzed by a hierarchy of
3 neural networks such as hierarchy 440 in Figure 4.

4

5 The neural networks filter the messages by looking at the intent of the
6 messages as indicated by the input vector generated in step 504. One embodiment of
7 the neural network hierarchy analyzes what words are used in a message, analyzes how
8 those words are used both independently and in relationship with each other, and con-
9 siders a classification for the message based on this analysis and on knowledge of other
10 messages.

11

12 As further description, in step 508, a first level neural network determines
13 if a message is more likely legitimate or junk. This step designates the message as
14 "likely good," which can include both good and bulk messages, and "likely spam,"
15 which can include bulk and spam messages.

16

17 In a step 509, likely good messages are analyzed by a second level neural
18 network to determine if they are good messages or bulk messages. Similarly, in a step
19 510, likely spam messages are analyzed by another second level neural network to de-
20 termine if they are bulk messages or spam messages.

21

1 In one embodiment, the neural networks “learn” (that is, are adjusted us-
2 ing known techniques for neural network learning, such as for example back-
3 propagation) from messages that are processed (and feedback from end-users in re-
4 sponse thereto), with the effect that the neural networks adapt to evolving anti-filtering
5 strategies that might be employed by senders of spam.

6

7 Good messages are handled at a step 511. These messages are sent to an
8 end-user recipient or to a mailbox for the end-user recipient.

9

10 Bulk messages are handled at a step 512. In one embodiment, bulk mes-
11 sages are tagged, for example by modifying their subject header, and sent to an end-
12 user recipient or to a mailbox for the end-user recipient. Alternatively, the spam filter
13 or spam filtering server can be configured by an administrator or user to treat bulk mes-
14 sages as good messages or as spam messages. This is indicated by the dashed lines in
15 Figure 5.

16

17 Spam messages are handled at a step 513. In one embodiment, these mes-
18 sages are blocked (that is, not sent to an end-user). The messages can be stored, for ex-
19 ample in a database, for later review and possibly retrieval in case of misidentification
20 of bulk or good messages as spam. In one embodiment, any misidentification (that is,
21 “false positives”) are used to further adjust (as described above) the neural networks in
22 order to help prevent similar misidentifications in the future.

1
2 In each of the steps 511, 512, and 513, the rules and the hierarchy of neural
3 networks can be dynamically updated and maintained based on the results of the classi-
4 fication and the characteristics (e.g., text and generated input vector) for the message.

5
6 Different categories and actions can be used in different embodiments of
7 the invention.

8
9 *Alternative Embodiments*

10
11 Although preferred embodiments are disclosed herein, many variations
12 are possible which remain within the concept, scope, and spirit of the invention. These
13 variations would become clear to those skilled in the art after perusal of this applica-
14 tion.

15
16 • Applications of the invention are not limited to embodiments in which only text
17 messages, or messages that are primarily text, are examined. The invention in-
18 cludes embodiments in which other data types, including pictures (such as for
19 example, still pictures or moving pictures) or sound (such as encoded sound),
20 code fragments (such as HTML, DHTML, Java, JavaScript, and the like), and
21 other elements of the message are considered when examining the message.

22

1 • Applications of the invention are not limited to embodiments in which only the
2 content of the headers or body of the message are examined. The invention in-
3 cludes embodiments in which other data about the message, including its time of
4 sending or receipt, its size, its method or path of transmission (such as for exam-
5 ple within the same enterprise or from a logically distant location), its protocol
6 for transmission (such as for example POP, IMAP, or SMTP mail, or variants
7 thereof), and whether that message includes attachments (and if so, their con-
8 tents, size, or data type), are examined in addition to, or instead of, aspects de-
9 scribed herein.

10

11 • Applications of the invention are not limited to embodiments in which only
12 email messages are examined. The invention includes embodiments in which
13 other types of messages, including web pages or files received using other proto-
14 cols (such as for example HTTP, HTTPS, FTP, or UDP) or messages of a com-
15 pletely different type (such as for example “instant messages,” messages using
16 peer-to-peer protocols, messages using HTTP tunneling, or application-specific
17 messages such as “NeoMail” on www.neopets.com), are examined. For exam-
18 ple, not intended to be limiting in any way, the invention might be configured to
19 block web pages or portions thereof, such as for example (1) advertising embed-
20 ded in web pages, (2) “pop-up” or “pop-under” web pages, or (3) web pages in-
21 cluding content inappropriate for children, inappropriate for a specific location

1 or workspace, offensive to a particular user, or simply unwanted by a particular
2 user.

3

4 • Applications of the invention are not limited to embodiments in which only hu-
5 man readable messages are examined. The invention includes embodiments in
6 which other types of messages, including CRM or ERP system messages, or other
7 forms of communication among and between multiple processors, are examined.
8 For example, not intended to be limiting in any way, the invention includes em-
9 bodiments in which bids or orders in an online auction or other transaction sys-
10 tem are examined for error, falsity, spoofing, or other factors (such as for exam-
11 ple, extreme market volatility).

12

13 • Applications of the invention are not limited to embodiments in which messages
14 are merely blocked or filtered. The invention includes embodiments in which
15 other types of actions, including forwarding those messages, ranking them for
16 priority, copying them to more permanent storage, dispatching them using
17 wireless communication (or related technologies, such as for example sending a
18 page or a telephone call). For example, the invention includes embodiments in
19 which it is determined whether to forward an entire message or to omit any at-
20 tachments, to a wireless recipient.

1 • Other applications of the invention include use to automatically prioritize mes-
2 sages for forwarding to wireless devices or for secondary review, to determine
3 which web pages to pre-load or cache, to detect malicious “pop-up” advertising
4 or web page JavaScript, to detect unwanted “instant messages” and other mes-
5 sage types, to automatically collect groups of messages into message digests, and
6 to automatically detect outgoing messages for virus content.

7

8 Those skilled in the art will recognize, after perusal of this application,
9 that these alternative embodiments are illustrative and in no way limiting.